Network Path State Machine

OIF Webinar

October 9, 2024

Eric Maniloff - Ciena

Overview

- Intro to CMIS
  - Data Path State Machine (DPSM):
    - Applications advertising - Key link to both DPSM & NPSM
    - DPSM Configuration: Commissioning and Provisioning
- The Network Path State Machine (NPSM)
  - What problems is this solving?
  - Implementation and operation
- Advanced Implementation
  - Multiplexing
  - Reconfiguration

CMIS: What & Why?

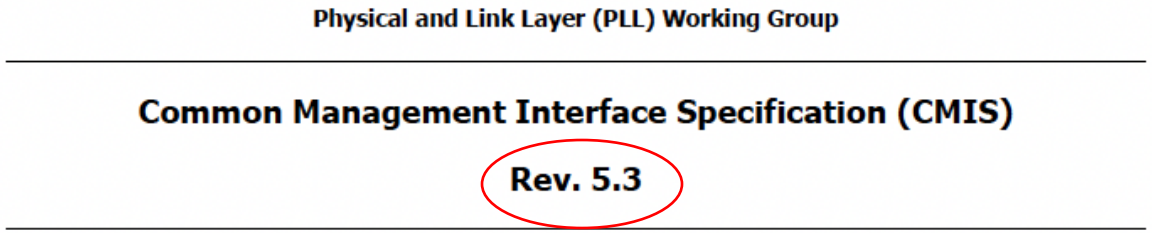

- CMIS in its own words:
  - *This **Common Management Interface Specification** (CMIS) defines a generic management communication interface together with a generic management interaction protocol between hosts and managed modules.*

- C = Common, with CMIS covering applications including QSFP-DD, OSFP, LPO, etc
  - CMIS is the base document, with supplementary documents covering specific applications

- Key People:
  - Editor: Stefan Langenbach (Cisco)
  - Management Track Chairs: Ian Alderdice (Ciena) & Gary Nicholl (Cisco)

- Where to find CMIS:
  - https://www.oiforum.com/technical-work/implementation-agreements-ias/#Management

CMIS Supplementary documents

- Extensions of CMIS cover the following application spaces

  - **CMIS-VCS** [6] defines signal integrity (SI) parameters that can be part of a so-called Versatile Control Set, generalizing and extending the original, fixed set of signal integrity controls described in section 6.2.3. Note that this functionality, when enabled, may "override" certain fields of the standard memory map, which, if applicable, will be specified in [6].
  - **C-CMIS** [7] provides specifications for modules with coherent transmission capabilities.
  - **CMIS-FF** [8] provides specifications of form factor specific management facilities typically implemented by additional and sometimes programmable discrete value hardware signals[2] that are available as per individual module form factor hardware specification[3].
  - **CMIS-LT** [9] provides specifications in support of host to module electrical link training. More precisely, CMIS-LT specifies both data structures and mechanisms to emulate a bidirectional message exchange between the electrical link endpoints of a host lane for the purpose of link training (LT). Note that the activities and behaviors of the acting link endpoints that reside in both host and module and that actively use those facilities are specified elsewhere.
  - **CMIS-ELSFP** [10] provides specifications for external laser source resource modules.

CMIS register space

- The register map is provided in CMIS 5.3 section 8
- CMIS pages are 128 Bytes each
- There is one Lower page (00h), and multiple Upper Pages (00h-FFh)
  - Pages are referenced in Hex, Addresses in Decimal – *Beware of confusion*
- CMIS defines Banked and Unbanked pages
- The register map provides full details of the implementation.
- Note: A shorthand for referencing addresses is:
  - Bytes: *Bank:Page:Byte ex: 0:10h:128-129*
  - Field: *Page:Byte.bits ex: 10h:128.7-4*

Module Terminology

(at the risk of stating the obvious)

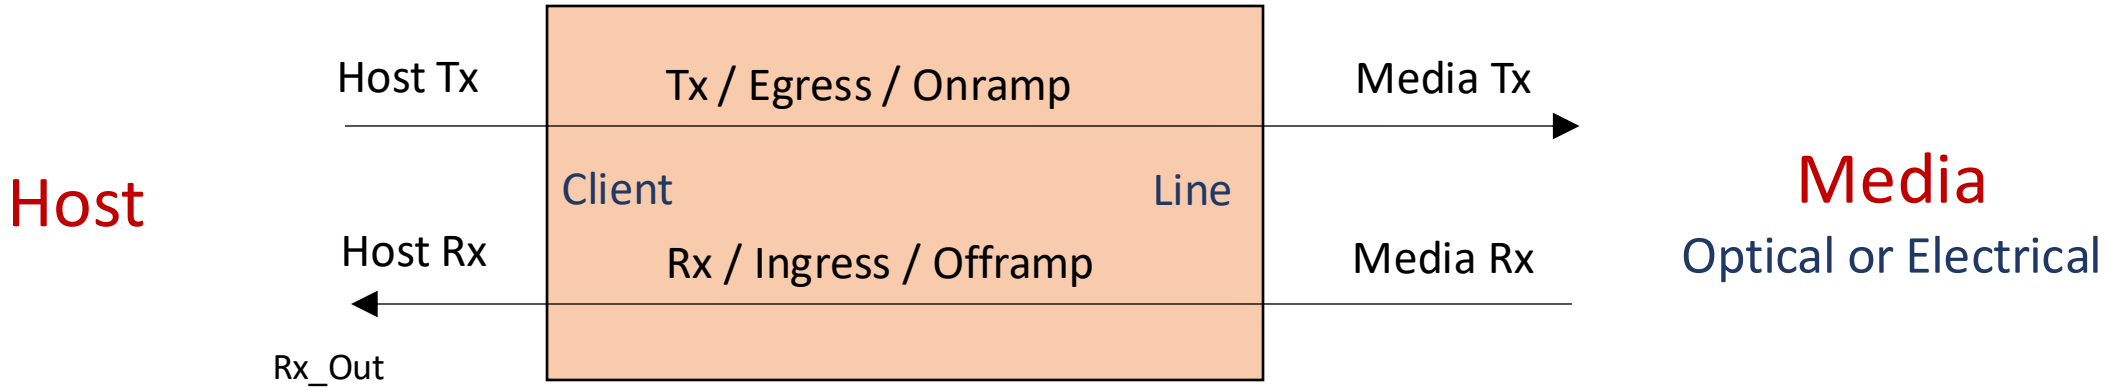

- Various naming is used in the different sources -
  - Tx/Rx is used throughout CMIS

CMIS Memory Map

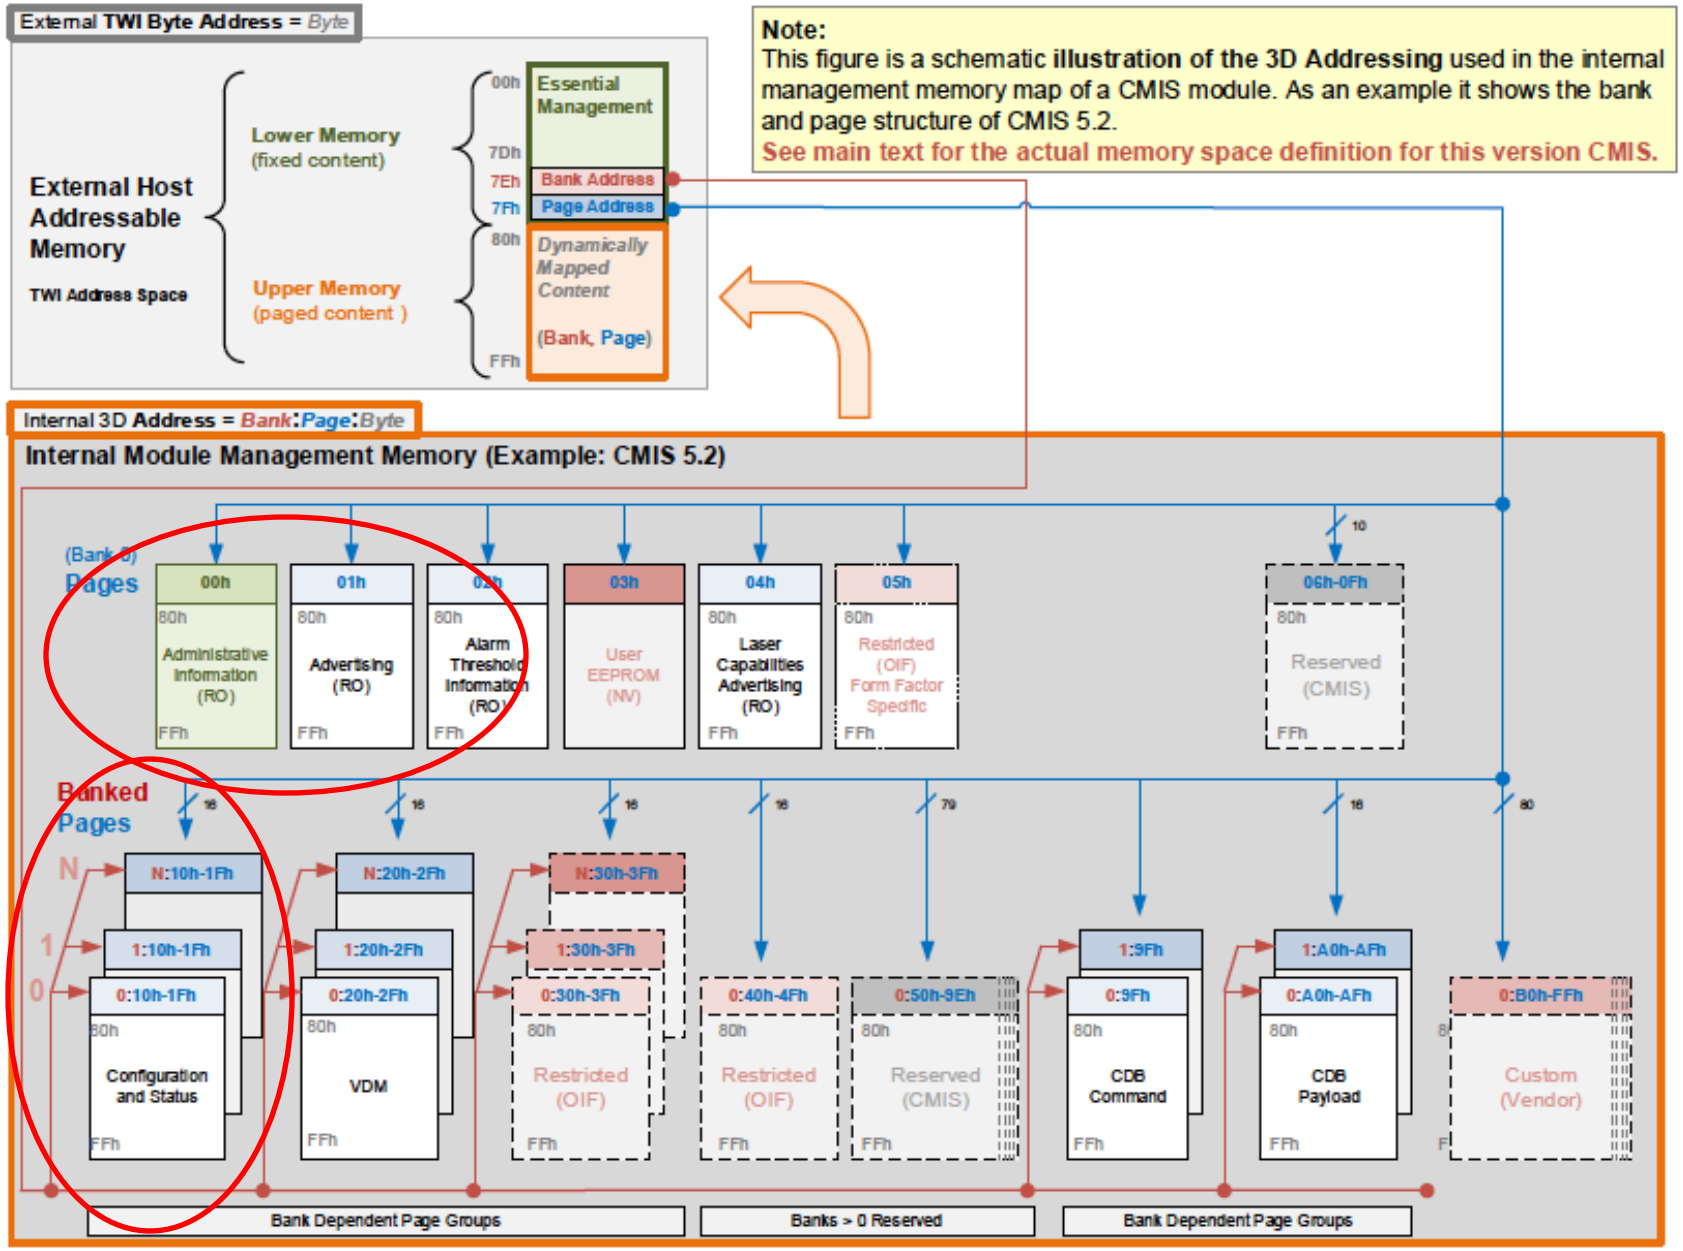

Figure 8-1 CMIS Module Memory Map (Conceptual View)

State Machines

- CMIS defines a number of Management State Machines:
  - Module State Machine governs the module behavior
  - Data Path State Machine (DPSM) is used to configure and manage the supported applications
    - In CMIS 5.0 the DPSM manages the full (Host + Media) configuration
    - In CMIS 5.1 the Network Path State Machine is added
  - The Network Path State Machine (NPSM) is used to configure and manage the media (line) side
    - Modules implementing the NPSM may have a mix of DP and NP applications
  - For hosts not implementing NPSM, DPSM is used for both host and media sides

Module State Machine

- The Module State Machine is illustrated here
- Timing of MgmtInit = 2 seconds
- Other durations advertised

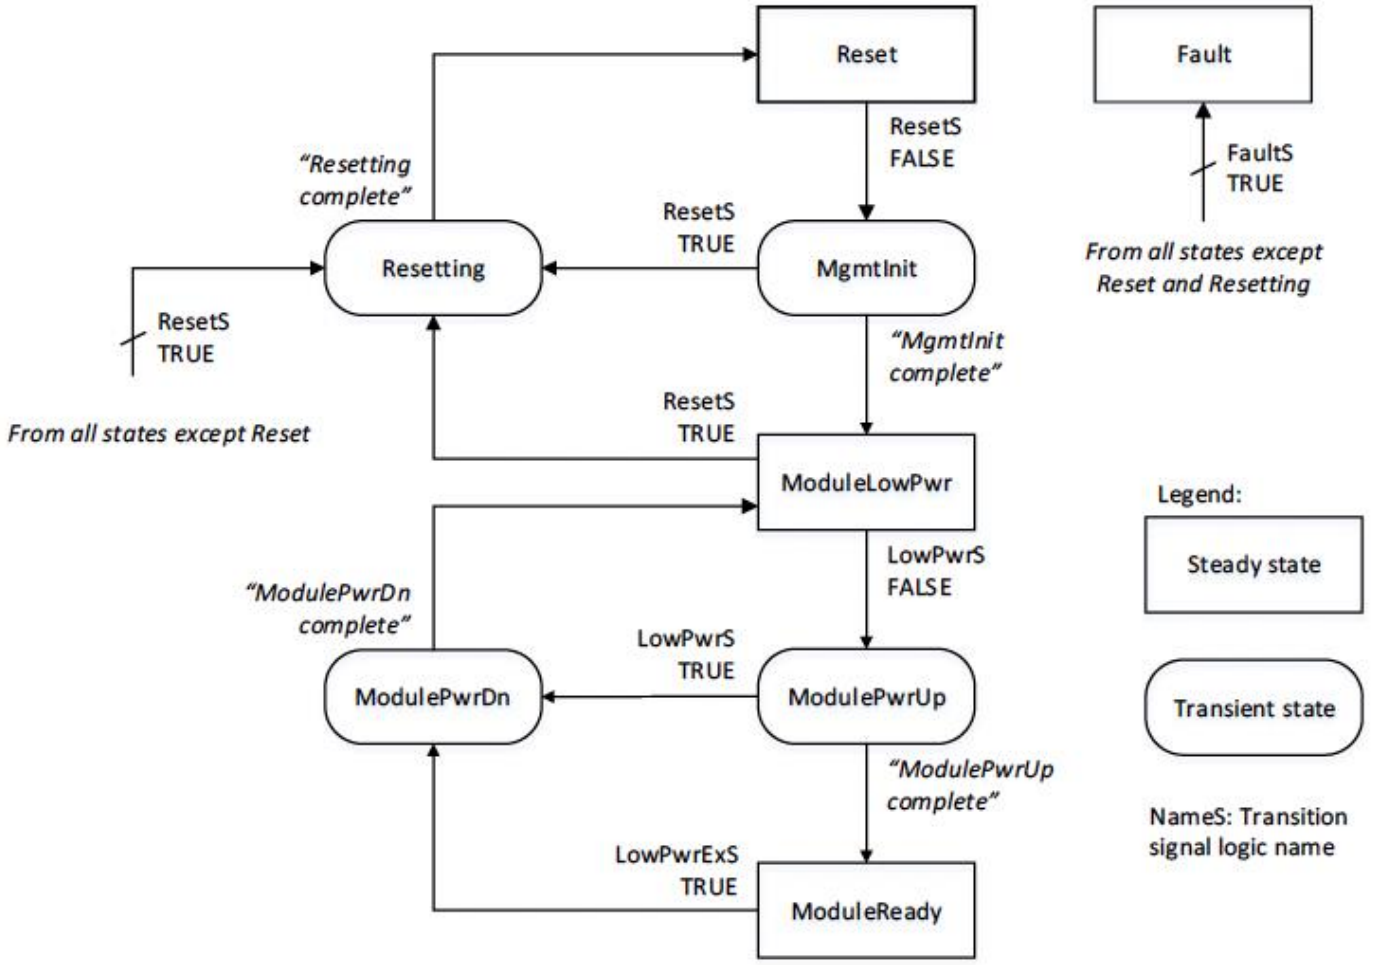

Figure 6-3 Paged Memory Module State Machine (MSM) State Transition Diagram

Configuring Data Path Applications

- CMIS defines a method of changing operating modes based on Application Descriptors
  - Application descriptors include both Media side (Line) and Host side (client) configurations.
  - A module will advertise one or more supported applications
    - Max of 15 Applications are supported in CMIS ≤ 5.2
- A host configures a Data Path by:
  - Selecting an application, data path settings, and corresponding host electrical settings in a Staged Control Set (SCS)
  - Loading the Staged Set into the Active Control Set (ACS)
- Applications are advertised in Lower Page and Upper Page 01h.
  - Extended applications are available in CMIS 5.3

Supported Applications Advertising

- Applications are advertised using Application Descriptors through appsels
  - Lower Page includes the first 8 Appsels
  - Upper Page 01h provides the next 7 Appsels

- CMIS 5.3 provides additional groups of 15 Normalized Applications Descriptors (NADs) advertised in page banked page 1Ch
  - NAD's are not currently covered in this document
  - From the standpoint of the DPSM and NPSM, NADs simply increase the number of applications that can be advertised

Application Definition Example: 400ZR / (400GbE + 4x100GbE)

Byte	Bits	AppSel Code	Name	Value	Description
85	7-0	N/A	Module Type encoding	02h	Optical Interfaces: SMF
86	7-0	001b	Host Electical Interface ID	11h	400GAUI-8 C2M
87	7-0		Module Media Interface ID	3Eh	400G-ZR
88	7-4		Host Lane Count	8h	8 Host Lanes
	3-0		Media Lane Count	1h	1 Media Lanes
89	7-0		Host Lane Assignment Options	01h	Permissable First Host Lane for Application: Lane 1
01h:176	7-0		Media Lane Assignment Options	01h	Permissible First Media Lane for Application: Lane 1
90	7-0	010b	Host Electical Interface ID	0Dh	100GAUI-2 C2M
91	7-0		Module Media Interface ID	3Eh	400G-ZR
92	7-4		Host Lane Count	2	
	3-0		Media Lane Count	1	
93	7-0		Host Lane Assignment Options	55h	1, 3, 5, 7
01h:177	7-0		Media Lane Assignment Options	01h	1 Media Lanes
94	7-0	011b	Host Electical Interface ID	FFh	End of List of Supported Applications
95	7-0		Module Media Interface ID	0	
96	7-4		Host Lane Count	0	
	3-0		Media Lane Count	0	
97	7-0		Host Lane Assignment Options	0	
01h:178	7-0		Media Lane Assignment Options	0	

SFF-8024 defines Interface ID's

DPSM Configuration

- Data Paths are configured by:
  1. Creating a Staged Set Application
  2. Applying the Staged Set
  3. Initializing the Application

Table 8-65 Staged Control Set 0, Data Path Configuration (Page 10h)

Byte	Bits	Field Name	Register Description	Type
145	7-4	AppSelCodeLane1	**SCS0::DPConfigLane1** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane1		
	0	ExplicitControlLane1		
146	7-4	AppSelCodeLane2	**SCS0::DPConfigLane2** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane2		
	0	ExplicitControlLane2		
147	7-4	AppSelCodeLane3	**SCS0::DPConfigLane3** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane3		
	0	ExplicitControlLane3		
148	7-4	AppSelCodeLane4	**SCS0::DPConfigLane4** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane4		
	0	ExplicitControlLane4		
149	7-4	AppSelCodeLane5	**SCS0::DPConfigLane5** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane5		
	0	ExplicitControlLane5		
150	7-4	AppSelCodeLane6	**SCS0::DPConfigLane6** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane6		
	0	ExplicitControlLane6		
151	7-4	AppSelCodeLane7	**SCS0::DPConfigLane7** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane7		
	0	ExplicitControlLane7		
152	7-4	AppSelCodeLane8	**SCS0::DPConfigLane8** See Table 8-64	RW Rqd.
	3-1	DataPathIDLane8		
	0	ExplicitControlLane8		

1

Table 8-63 Staged Control Set 0, Apply Triggers (Page 10h)

Byte	Bits	Field Name	Register Description	Type
143	7	ApplyDPInitLane8	**SCS0::ApplyDPInitLane<i>** 0b: No action for host lane <i> 1b: Trigger the **Provision** procedure using the Staged Control Set **0** settings for host lane <i>, with feedback provided in the associated ConfigStatusLane<i> field Restriction: This byte must be written in a single-byte WRITE *Note: See Table 6-3 for preconditions and state dependencies*	WO Rqd.
	6	ApplyDPInitLane7		
	5	ApplyDPInitLane6		
	4	ApplyDPInitLane5		
	3	ApplyDPInitLane4		
	2	ApplyDPInitLane3		
	1	ApplyDPInitLane2		
	0	ApplyDPInitLane1		

2

Table 8-61 Data Path initialization control (Page 10h:128)

Byte	Bit	Field Name	Register Description	Type
128	7	DPDeinitLane8	**DPDeinitLane<i>** Data Path initialization control for host lane <i> 0b: Initialize the Data Path associated with host lane 1b: Deinitialize the Data Path associated with host lane All lanes of a Data Path must have the same value *Note: These bits represent static requests, not trigger events*	RW Rqd.
	6	DPDeinitLane7		
	5	DPDeinitLane6		
	4	DPDeinitLane5		
	3	DPDeinitLane4		
	2	DPDeinitLane3		
	1	DPDeinitLane2		
	0	DPDeinitLane1		

3

Data Path State Machine

- The Data Path State Machine is illustrated here
- CMIS defaults DPDeactivated to false
- For DP applications this state machine controls the line and host sides
- In CMIS 5.1 the NPSM was introduced to handle multiplexing/uniplexing applications

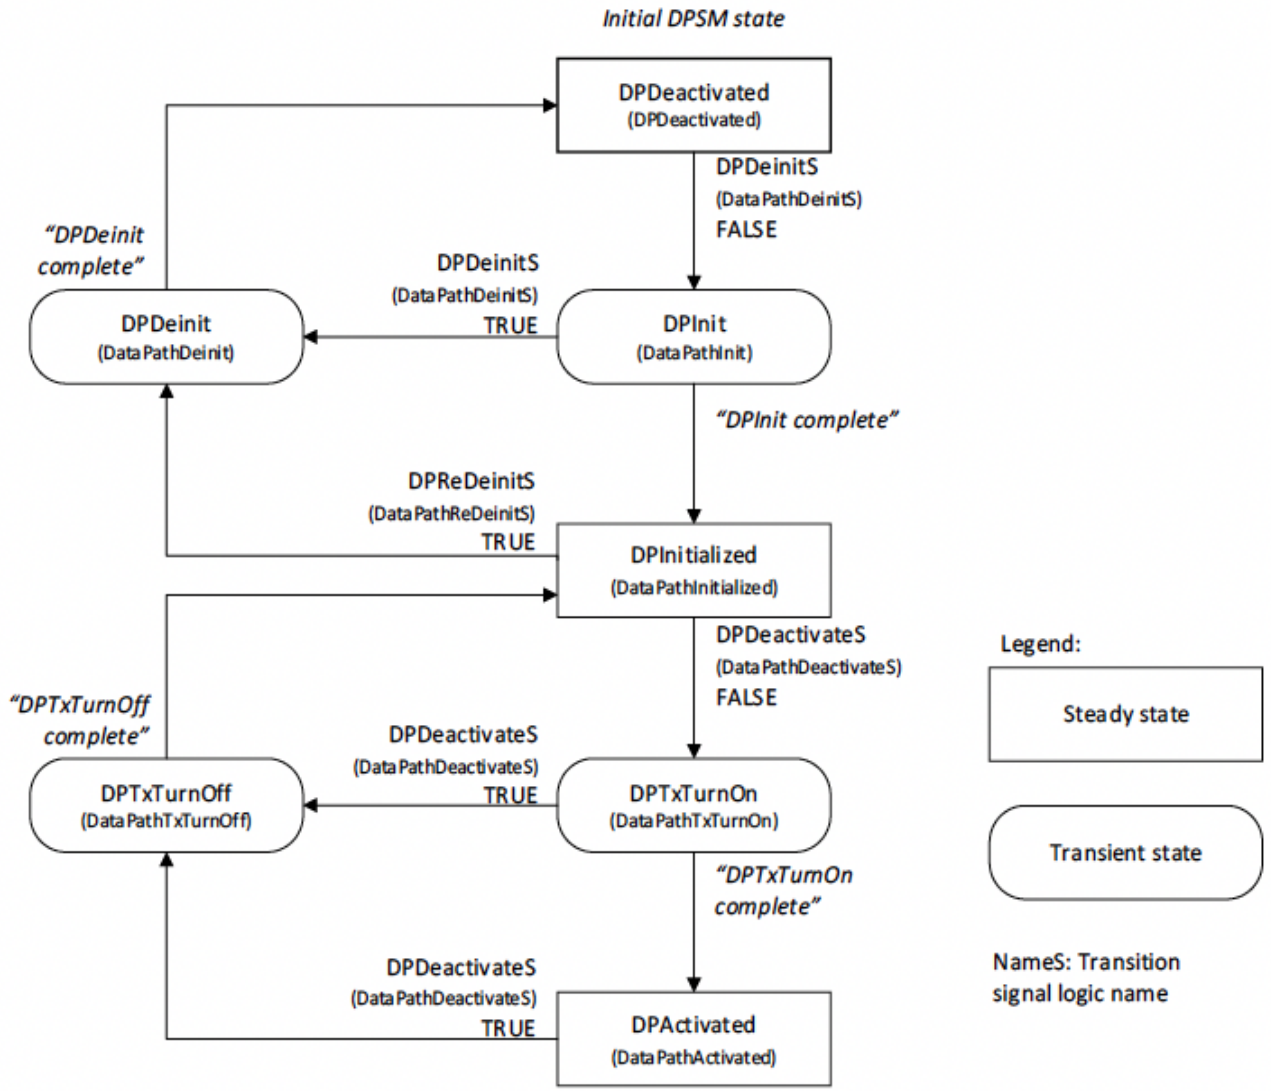

Figure 6-5 Data Path State Machine (DPSM) State Transition Diagram

DP State and Active Set

- Page 11h displays the DP state, as well as the Active Contol Set

Table 8-83 Lane-associated Data Path States (Page 11h)

Byte	Bit	Field Name	Register Description (DPStateHostLane<i>)	Type
128	7-4	DPStateHostLane2	Data Path State of host lane 2 (see Table 8-84)	RO
	3-0	DPStateHostLane1	Data Path State of host lane 1 (see Table 8-84)	Rqd.
129	7-4	DPStateHostLane4	Data Path State of host lane 4 (see Table 8-84)	RO
	3-0	DPStateHostLane3	Data Path State of host lane 3 (see Table 8-84)	Rqd.
130	7-4	DPStateHostLane6	Data Path State of host lane 6 (see Table 8-84)	RO
	3-0	DPStateHostLane5	Data Path State of host lane 5 (see Table 8-84)	Rqd.
131	7-4	DPStateHostLane8	Data Path State of host lane 8 (see Table 8-84)	RO
	3-0	DPStateHostLane7	Data Path State of host lane 7 (see Table 8-84)	Rqd.

Table 8-84 Data Path State Encoding

Encoding	State
0h	**Reserved**
1h	DPDeactivated (or unused lane)
2h	DPInit
3h	DPDeinit
4h	DPActivated
5h	DPTxTurnOn
6h	DPTxTurnOff
7h	DPInitialized
8h-Fh	**Reserved**

Table 8-93 Active Control Set, Provisioned Data Path Configuration (Page 11h)

Byte	Bits	Field Name	Register Description	Type
206	7-4	AppSelCodeLane1	**ACS::DPConfigLane1**	RO
	3-1	DataPathIDLane1	See Table 8-92	Rqd.
	0	ExplicitControlLane1		
207	7-4	AppSelCodeLane2	**ACS::DPConfigLane2**	RO
	3-1	DataPathIDLane2	See Table 8-92	Rqd.
	0	ExplicitControlLane2		
208	7-4	AppSelCodeLane3	**ACS::DPConfigLane3**	RO
	3-1	DataPathIDLane3	See Table 8-92	Rqd.

⋮

DP Configuration status

- Error Codes:
  - 11h:202-205

Table 8-91 Configuration Command Execution and Result Status Codes (Page 11h)

Encoding	Name	Value Description
0h	ConfigUndefined	No status information available (initial register value)
1h	ConfigSuccess	**Positive Result Status:** The last accepted configuration command has been completed successfully
2h	ConfigRejected	**Negative Result Status (2h-Bh. Dh-Fh):** Configuration rejected: unspecific validation failure
3h	ConfigRejectedInvalidAppSel	Configuration rejected: invalid AppSel code
4h	ConfigRejectedInvalidDataPath	Configuration rejected: invalid set of lanes for AppSel
5h	ConfigRejectedInvalidSI	Configuration rejected: invalid SI control settings
6h	ConfigRejectedLanesInUse	Configuration rejected: some lanes not in DPDeactivated
7h	ConfigRejectedPartialDataPath	Configuration rejected: lanes are only subset of DataPath
8h	-	**Reserved** (other validation failures)
9h	-	
Ah	-	
Bh	–	
Ch	ConfigInProgress	**Execution Status:** A configuration command is still being processed by the module; a new configuration command is **ignored** for this lane while ConfigInProgress.
Dh	-	**Custom** Configuration rejected for custom reasons
Eh	-	
Fh	-	

Network Path State Machine

- The Network Path State machine is defined to allow the Media side to be brought up independent of the host side

- This allows:
  - Multiplexing applications in which some (or all) ports are left un-initialized during the line side initialization
  - Applications in which the host traffic is not available when the line is initialized
  - Host traffic reconfiguration

- Multiplexed clients can be provisioned without the NPSM, by provisioning all host lanes in a Staged Set, and initializing simultaneously
  - This requires valid electrical signals on all host lanes when the data paths are initialized.

- Selection between mission-mode traffic and module generated replacement signals is supported

Network Path (NP) module – conceptual block diagram

In a NP module one or more data paths (where each data path is associated with a separate host interface) are mapped onto the media interface associated with a network path

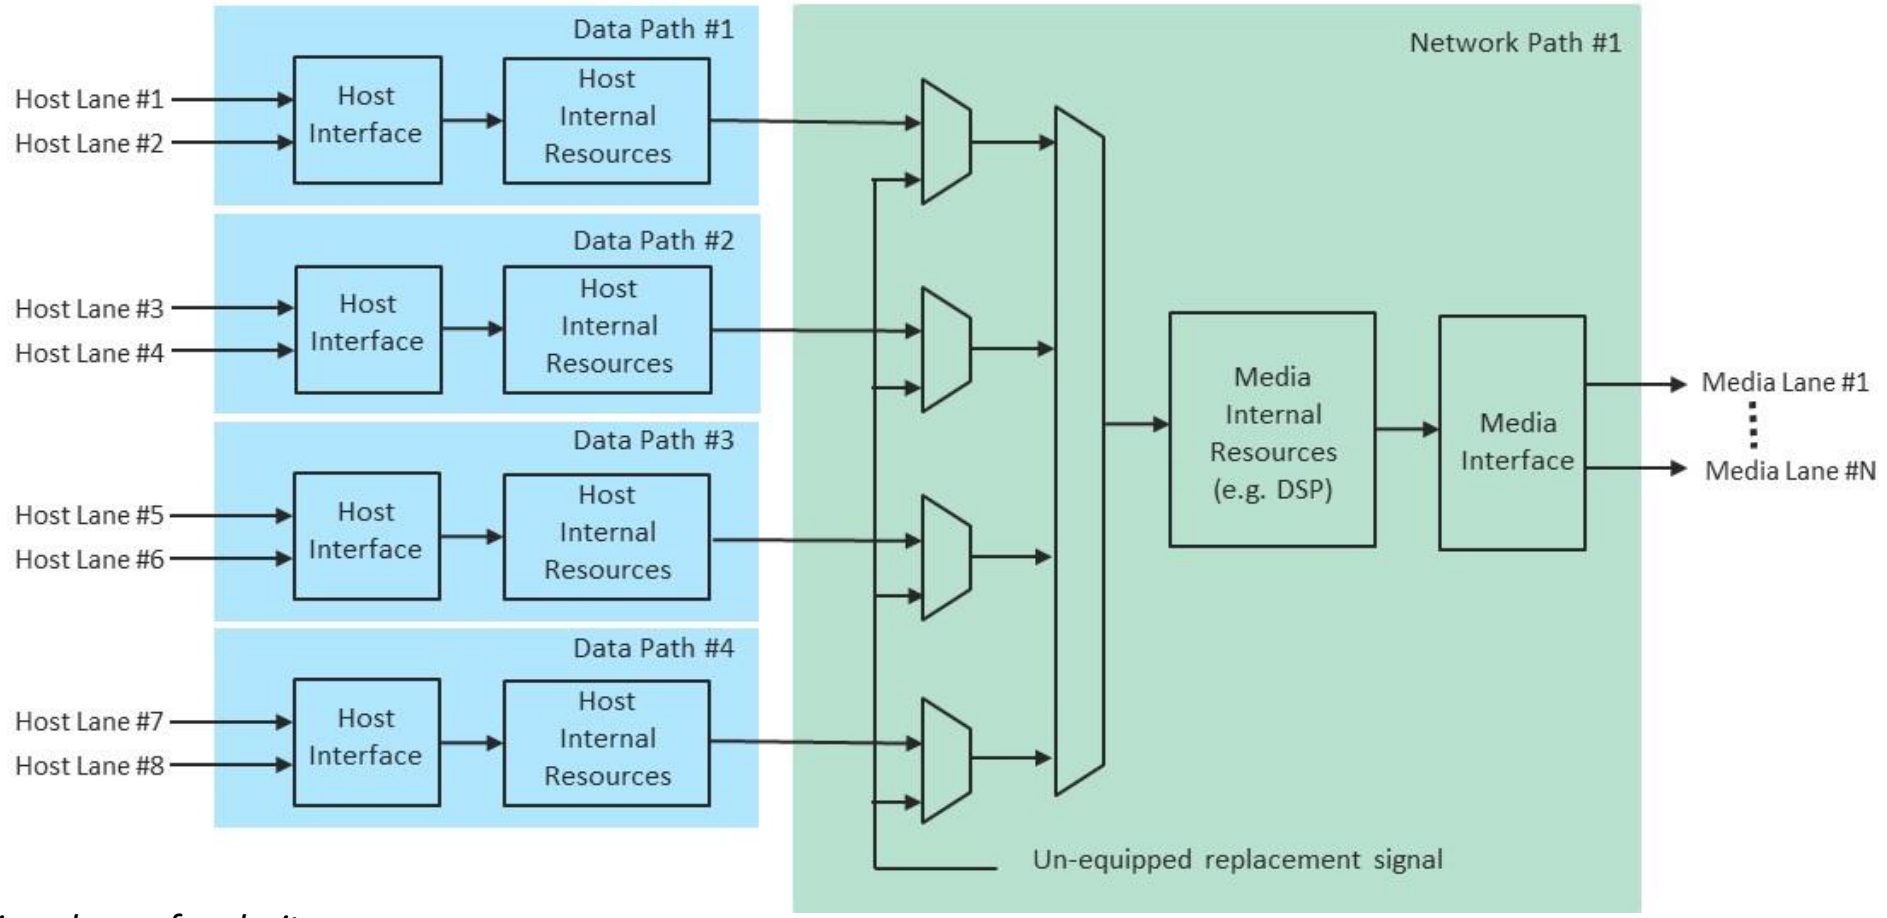

Note: Only one direction shown for clarity.

Application Overview

DP Application

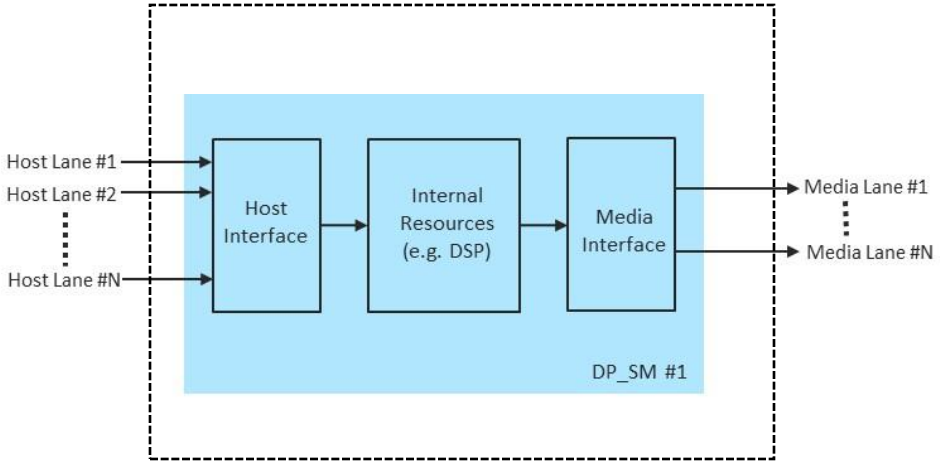

NP Application

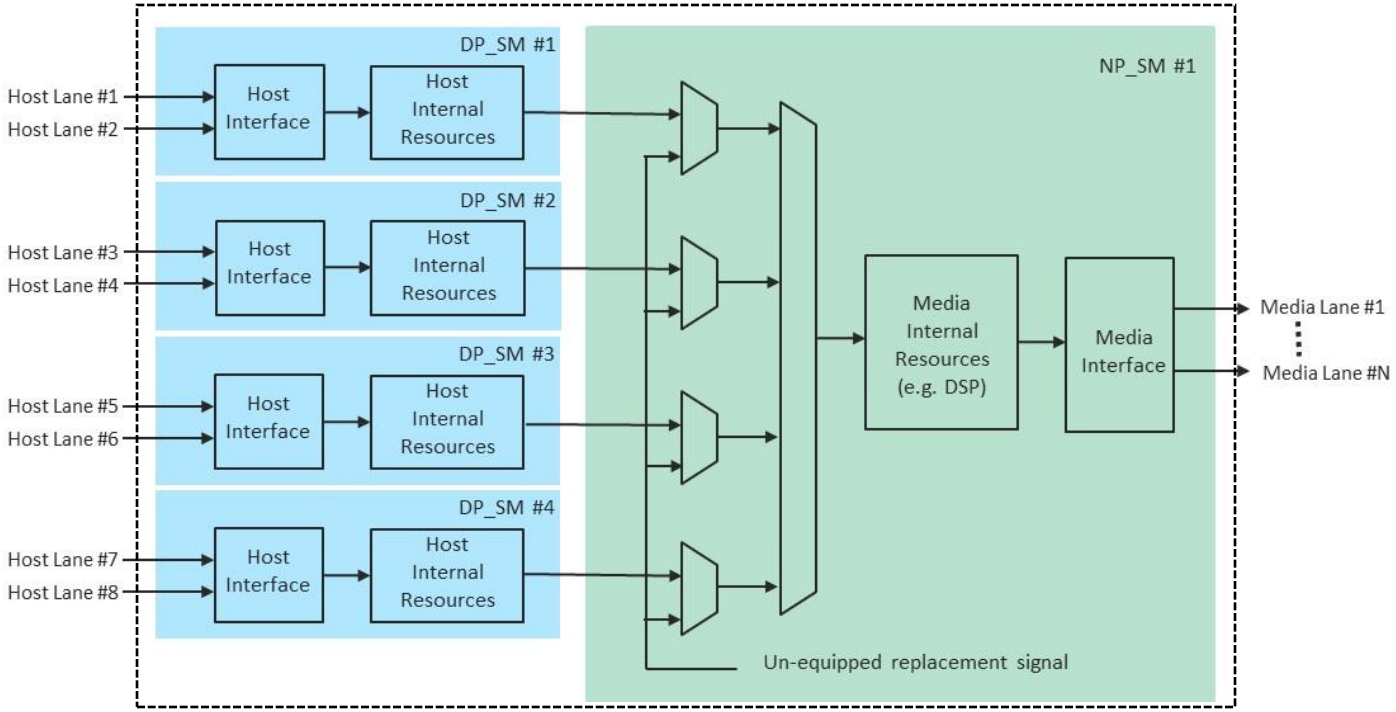

- In both cases an Application associates a group of host lanes to a group of media lanes
- A DP Application contains a single data path
- A NP Application contains a single network path and one or more host paths

NPSM Provisioning

- Applications advertised in Lower Page, Upper Page 01h are used to define host and media identifiers – This is unchanged from the DPSM
  - New advertising bits identify an application as a DP or NP application

- For NP applications:
  - The Media side provisioning is performed through the NPSM
  - The **Host Paths** are controlled by a modified DPSM
    - **DPDeactivateS** is **forced to TRUE**, limiting the steady states achievable to DPDeactivated and DPInitialized

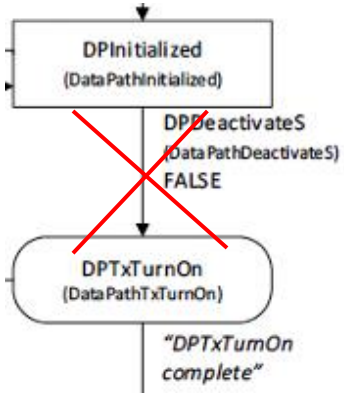

- Application Selectors are loaded using the Page 10h Staged Sets

- The Network Path is configured using the Network Path Configuration Staged Sets in Page 16h

- NPSM instances will default to values corresponding to the default application:

> All **NPSM** instances required to represent the power-up default Application defined in the Network Path Configuration field values of the Active Control Set are initially created and set-up during the **MgmtInit** state.

CMIS NPSM

NPSM largely mirrors the DPSM

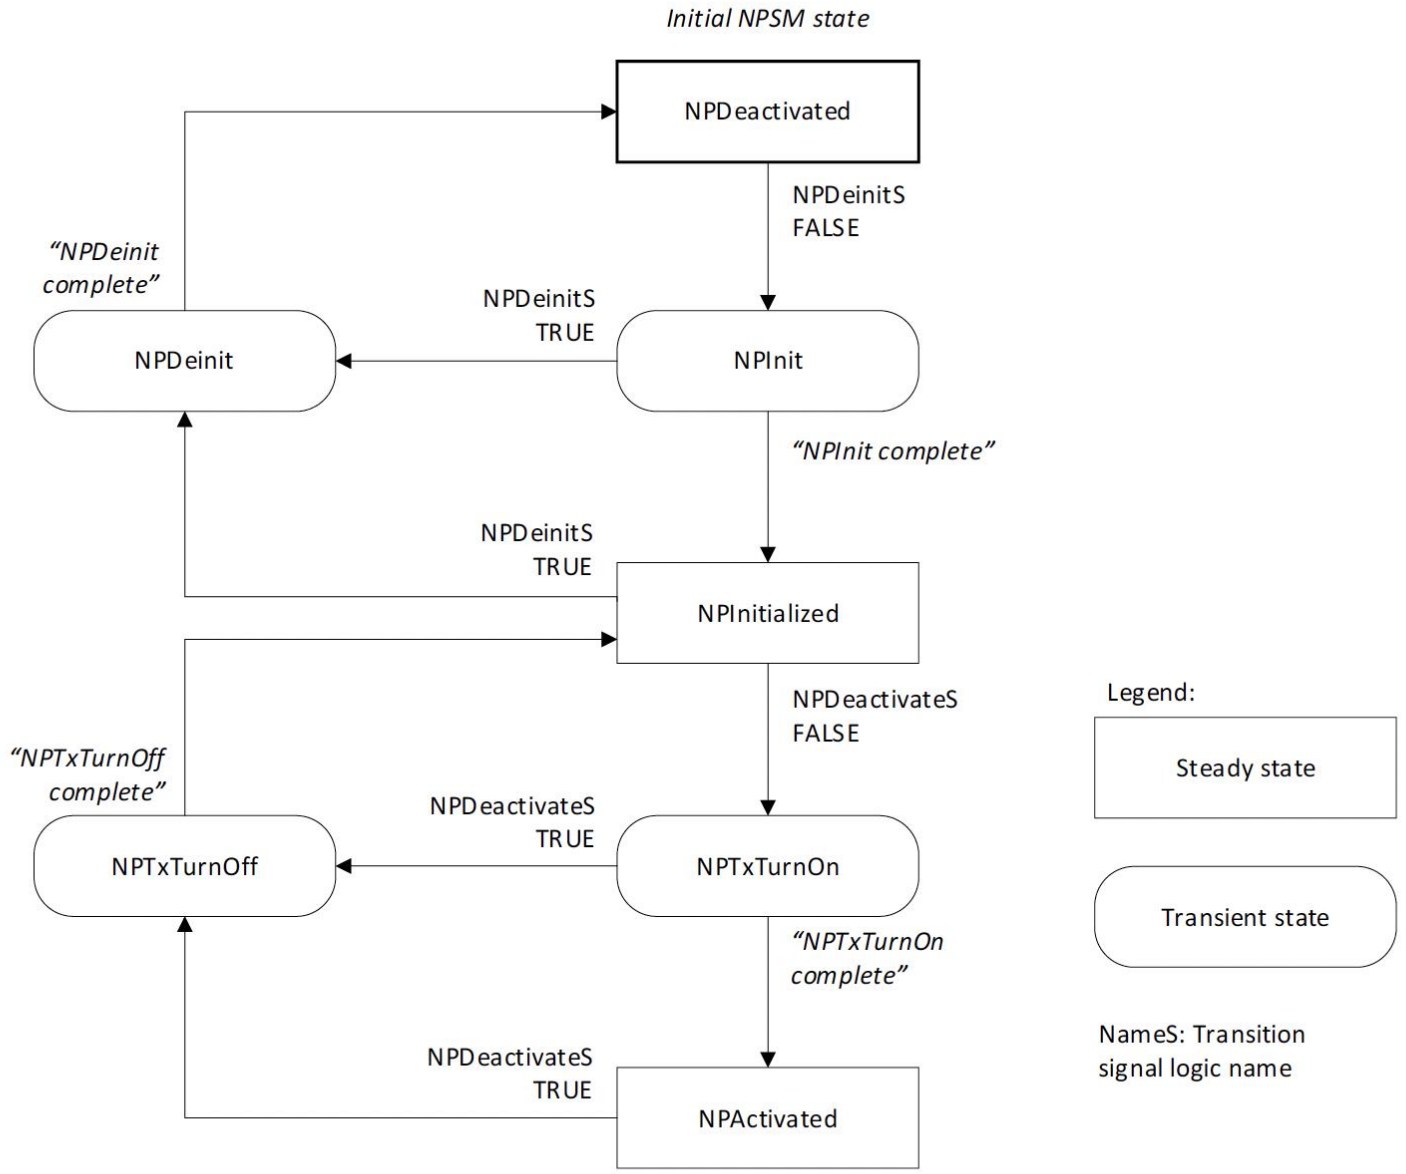

Figure 7-6 Network Path State Machine (NPSM) State Transition Diagram

Network Path Applications

- For a Network Path (NP) application the host is still controlled by the DPSM
  - In a NP application the Host Path is controlled by the legacy DP registers
  - When operating as a HP, the DP *only* operates on the host portion of the application
  - Achievable steady states are DPDeactivated and DPInitialized
- Media side and Tx control is performed by the NPSM
- In NP Applications the HP and NP are independently provisioned and initialized

HPs and NP are then (requested to be) **provisioned** into the relevant Active Control Set or NP Active Control Set in separate steps, by writing to the associated ApplyDPInit trigger register or ApplyNPInit register, respectively. As usual, the provisioning procedure includes (partial) validation prior to copying a Staged Control Set into the associated Active Control Set.

Page 16h: Network Path Functionality

Table 8-124 Page 16h Overview

Byte	Size (bytes)	Subject Area	Description
128-159	**32**	**Provisioning**	
128-135	8	NP Staged Control Set 0	Lane to NP Assignments Provisioning – Staged Control Set 0
136-143	8	NP Staged Control Set 1	Lane to NP Assignments Provisioning – Staged Control Set 1
144-159	16	-	**Reserved[16]**
160-175	**16**	**Control**	
160	1	NP Control	Network Path initialization control
161	1	-	**Reserved[1]**
162-163	2	NP Source Selectors	Signal source selection at the Network Path connection points
164-175	12	-	**Reserved[12]**
176-191	**16**	**Command & Response**	
176	1	NP Apply SCS 0	Apply command for NP Staged Control Set 0
177	1	NP Apply SCS 1	Apply command for NP Staged Control Set 1
178-181	4	Configuration Status	Status of most recent Network Path configuration command
182-191	10	-	**Reserved[10]**
192-223	**32**	**Status**	
192-199	8	NP Active Control Set	Provisioned Network Path Configuration
200-203	4	Network Path Status	Network Path State Machine state of each NP media lane
204	1	NPInitPending Condition	Commissioning status (NPInitPending condition)
205-223	19	-	**Reserved[19]**
224-255	**32**	**Advertisement**	
224-225	2	NPSM Max Durations	Maximum durations for all NPSM transient states
226	1	Options	Miscellaneous options
227	1	-	**Reserved[1]**
228-247	20	Mixed Multiplex Support	Advertising for mixed HP multiplexing support
248-249	2	Application Advertisement	Application Advertisement Extensions
250-255	6	-	**Reserved[4]**

Note: NP State Change Flags are in Page 17h

NPSM Applications Advertisement

- Registers 16h:248-249 define each application as being a DP Application or NP Application
- If Bytes 248 & 249 = 00h, the applications are all managed entirely by the DPSM
  - Modules not implementing NPSM (including pre CMIS 5.1) will set these bits to zero
  - Hosts not implementing the NPSM (CMIS 5.0) will ignore these bits
  - A mix of DP and NP applications is possible

Byte	Bits	Field Name	Register Description	Type
248	7	ExtAppDescriptor15	**ExtAppDescriptor<i>**	RW
	6	ExtAppDescriptor14	The Application Descriptor identified by AppSel<i>:	Rqd.
	5	ExtAppDescriptor13	0b: describes a supported DP Application	
	4	ExtAppDescriptor12	1b: (partially) describes a supported NP Application	
	3	ExtAppDescriptor11		
	2	ExtAppDescriptor10	*Note: The value is irrelevant when the Application Descriptor identified by AppSel<i> is unused.*	
	1	ExtAppDescriptor9		
	0	ExtAppDescriptor8		
249	7	ExtAppDescriptor7		
	6	ExtAppDescriptor6		
	5	ExtAppDescriptor5		
	4	ExtAppDescriptor4		
	3	ExtAppDescriptor3		
	2	ExtAppDescriptor2		
	1	ExtAppDescriptor1		
	0	-		

Application Definition Example: 400ZR (400GbE & 4x100GbE): NP Application

Byte	Bits	AppSel Code	Name	Value	
85	7-0	N/A	Module Type encoding	02h	
86	7-0	001b	Host Electical Interface ID	11h	
87	7-0		Module Media Interface ID	3Eh	
	7-4		Host Lane Count	8h	
88	3-0		Media Lane Count	1h	
89	7-0		Host Lane Assignment Options	01h	
01h:176	7-0		Media Lane Assignment Options	01h	
90	7-0	010b	Host Electical Interface ID	0Dh	
91	7-0		Module Media Interface ID	3Eh	
	7-4		Host Lane Count	2	
92	3-0		Media Lane Count	1	
93	7-0		Host Lane Assignment Options	55h	
01h:177	7-0		Media Lane Assignment Options	01h	
94	7-0	011b	Host Electical Interface ID	FFh	
95	7-0		Module Media Interface ID	0	
	7-4		Host Lane Count	0	
96	3-0		Media Lane Count	0	
97	7-0		Host Lane Assignment Options	0	
01h:178	7-0		Media Lane Assignment Options	0	

Page 16h

Byte	Bit	Name	Value Hex	Value	App Type
248	7	ExtAppDescriptor15	00h	0	-
	6	ExtAppDescriptor14		0	-
	5	ExtAppDescriptor13		0	-
	4	ExtAppDescriptor12		0	-
	3	ExtAppDescriptor11		0	-
	2	ExtAppDescriptor10		0	-
	1	ExtAppDescriptor9		0	-
	0	ExtAppDescriptor8		0	-
249	7	ExtAppDescriptor7	06h	0	-
	6	ExtAppDescriptor6		0	-
	5	ExtAppDescriptor5		0	-
	4	ExtAppDescriptor4		0	-
	3	ExtAppDescriptor3		0	-
	2	ExtAppDescriptor2		1	NP
	1	ExtAppDescriptor1		1	NP
	0	-		0	Unused

Legacy modules would default Page 16h to 0's --> DP Only

NP Configuration – high level view of a possible stepped configuration

- In Module Low Power:
  - Host sets DPSM to DPDeinit (Page 10h:128) and NPSM to NPDeinit (16h:160)
- Application is configured into the Active Set (p10h), copying from a staged set, as in DP applications
- NP Staged Control Sets 0 &/or 1 (16h:128-135, 16h:136-143) are configured
- Active NP is written from NP SS 0 (16h:176) or NP SS 1 (16h:177) to the Active Set (16h:192-199) using ApplyNPInit
  - ApplyNPInit does not Initialize the NP
  - Configuration Status is available in 16h:178-181
  - NPInitPending is advertised in 16h:204 when the Active Set has not been initialized
- Signal Source selection on the Rx and Tx directions configured in 16h:162-163
- Network Path is initialized by setting 16h:160.n bit n to 0 for **all** host lanes feeding the NP, removing Deinit
- Host paths are initialized by setting 10h:128.n bit n to 0 for all host lanes being used
  - Host and Network initialization can be performed in either order

Network Path Provisioning

- The Network Path is provisioned to indicate the lanes feeding the NP:
  - NPID Indicates the first host lane feeding the NP
  - Each lane indicates whether it is in use as part of an NP application

Table 8-133 Network Path Provisioning per Lane (NPConfigLane<i>)

Lane	Bits	Field Name	Register Description	Type
<i>	7-4	-	**Reserved**	RW
	3-1	NPID	**SCS<k>::NPIDLane<i>** If host lane <i> feeds the Network Path of an NP Application instance, the NPIDLane<i> field stores the **Network Path ID** of that Network Path, which is defined as the number of the first host lane feeding the Network Path, decremented by one. If lane<i> is unused (NPInUseLane<i> = 0), the value of NPIDLane<i> is ignored. *Note: All lanes of the Network Path of an Application that spans multiple host lanes have the same NPID.* *Note: For example, the NPID of a Network Path carrying the HP of host lane 1 is 0 and the NPID of a NP where host lane 5 is the lowest lane number feeding the NP is 4.*	Rqd.
	0	NPInUse	**SCS<k>::NPInUseLane<i>** 0b: host lane <i> is either part of the Data Path of a DP Application or it is unused 1b: host lane <i> is part of a Host Path that feeds the Network Path identified by the NPID field in a NP Application	

Staged NP Sets

- Two staged sets are available.
  - NP Stage sets configure which host lanes are part of an NP

Table 8-134 Staged Control Set 0, Network Path Configuration (Page 16h)

Byte	Bits	Field Name	Register Description	Type
128	7-4	-	**SCS0::NPConfigLane1** See Table 8-133	RW Rqd.
	3-1	NPIDLane1		
	0	NPInUseLane1		
129	7-4	-	**SCS0::NPConfigLane2** See Table 8-133	RW Rqd.
	3-1	NPIDLane2		
	0	NPInUseLane2		
130	7-4	-	**SCS0::NPConfigLane3** See Table 8-133	RW Rqd.
	3-1	NPIDLane3		
	0	NPInUseLane3		
131	7-4	-	**SCS0::NPConfigLane4** See Table 8-133	RW Rqd.
	3-1	NPIDLane4		
	0	NPInUseLane4		
132	7-4	-	**SCS0::NPConfigLane5** See Table 8-133	RW Rqd.
	3-1	NPIDLane5		
	0	NPInUseLane5		
133	7-4	-	**SCS0::NPConfigLane6** See Table 8-133	RW Rqd.
	3-1	NPIDLane6		
	0	NPInUseLane6		
134	7-4	-	**SCS0::NPConfigLane7** See Table 8-133	RW Rqd.
	3-1	NPIDLane7		
	0	NPInUseLane7		
135	7-4	-	**SCS0::NPConfigLane8** See Table 8-133	RW Rqd.
	3-1	NPIDLane8		
	0	NPInUseLane8		

NPIDLane<i> is defined as the number of the first host lane feeding the Network Path, decremented by one.

NP SCS examples

Address	800GAUI-8	2x400GAUI-4	400GAUI-4	400GAUI-4
128	0001b	0001b	0001b	0000b
129	0001b	0001b	0001b	0000b
130	0001b	0001b	0001b	0000b
131	0001b	0001b	0001b	0000b
132	0001b	0001b	0000b	1001b
133	0001b	0001b	0000b	1001b
134	0001b	0001b	0000b	1001b
135	0001b	0001b	0000b	1001b
		800G media side		400G media side

NP Control

- ApplyNPInit and NPDeinit are used to control NP configuration and the NPSM

Table 8-130 Staged Control Set 0, Apply Triggers (Page 16h)

Byte	Bits	Field Name	Register Description	Type
176	7	ApplyNPInitLane8	**SCS0::ApplyNPInitLane<i>** 0b: No action for host lane <i> 1b: Trigger the **Provision** procedure using the NP Staged Control Set **0** settings for host lane <i>, with feedback provided in the associated NPConfigStatusLane<i> field Restriction: This byte must be written in a single-byte WRITE	WO Rqd.
	6	ApplyNPInitLane7		
	5	ApplyNPInitLane6		
	4	ApplyNPInitLane5		
	3	ApplyNPInitLane4		
	2	ApplyNPInitLane3		
	1	ApplyNPInitLane2		
	0	ApplyNPInitLane1		

Table 8-131 Staged Control Set 1, Apply Triggers (Page 16h)

Byte	Bits	Field Name	Register Description	Type
177	7	ApplyNPInitLane8	**SCS1::ApplyNPInitLane<i>** 0b: No action for host lane <i> 1b: Trigger the **Provision** procedure using the NP Staged Control Set **0** settings for host lane <i>, with feedback provided in the associated NPConfigStatusLane<i> field Restriction: This byte must be written in a single-byte WRITE	WO Rqd.
	6	ApplyNPInitLane7		
	5	ApplyNPInitLane6		
	4	ApplyNPInitLane5		
	3	ApplyNPInitLane4		
	2	ApplyNPInitLane3		
	1	ApplyNPInitLane2		
	0	ApplyNPInitLane1		

Table 8-128 Network Path Initialization Control (Page 16h)

Byte	Bits	Field Name	Register Description	Type
160	7	NPDeinitLane8	**NPDeinitLane<i>** Initialization control for the Network Path fed by host lane <i> 0b: Initialize the Network Path associated with host lane <i> 1b: Deinitialize the Network Path associated with host lane <i> All host lanes feeding one Network Path must have the same NPDeinitLane<i> value set *Note: These bits represent static requests, not trigger events*	RW Rqd.
	6	NPDeinitLane7		
	5	NPDeinitLane6		
	4	NPDeinitLane5		
	3	NPDeinitLane4		
	2	NPDeinitLane3		
	1	NPDeinitLane2		
	0	NPDeinitLane1		

NP Configuration Status

- On an apply, the configuration is evaluated, and a status is returned
  - Note: Apply will be rejected unless the NP is in the NPdeactivated State

Table 8-132 NP Configuration Command Status registers (Page 16h)

Byte	Bit	Field Name	Field Description	Type
178	7-4	NPConfigStatusLane2	**NPConfigStatusLane<i>** Provisioning Command Execution / Result Status for the Network Path of host lane <i>, during and after the most recent configuration command. See Table 8-133 for the encoding of values. *Note: There is no feedback to the host when an Apply trigger is ignored after failed readiness test (when another configuration is still in progress)*	RO Rqd.
	3-0	NPConfigStatusLane1		
179	7-4	NPConfigStatusLane4		RO Rqd.
	3-0	NPConfigStatusLane3		
180	7-4	NPConfigStatusLane6		RO Rqd.
	3-0	NPConfigStatusLane5		
181	7-4	NPConfigStatusLane8		RO Rqd.
	3-0	NPConfigStatusLane7		

Table 8-133 NP Configuration Command Execution and Result Status Codes (Page 16h)

Encoding	Name	Value Description
0h	ConfigUndefined	No status information available (initial register value)
1h	ConfigSuccess	**Positive Result Status**: The last accepted configuration command has been completed successfully
2h	ConfigRejected	**Negative Result Status (2h-Bh. Dh-Fh)**: Configuration rejected: unspecific validation failure
3h	ConfigRejectedInvalidAppSel	Configuration rejected: invalid AppSel codes
4h	ConfigRejectedInvalidNetworkPath	Configuration rejected: invalid set of lanes for AppSel
5h	-	**Reserved**
6h	ConfigRejectedLanesInUse	Configuration rejected: some lanes not in NPDeactivated
7h	ConfigRejectedPartialNetworkPath	Configuration rejected: lanes are only subset of Network Path
8h	-	**Reserved** (other validation failures)
9h	-	
Ah	-	
Bh	-	
Ch	ConfigInProgress	**Execution Status**: A configuration command is still being processed by the module; a new configuration command is **ignored** for this lane while ConfigInProgress.
Dh	-	**Custom** Configuration rejected for custom reasons
Eh	-	
Fh	-	

Network Path Active Control Set

- The NP Active Control Set displays the configured NP

Table 8-134 NP Active Control Set, Network Path Configuration (Page 16h)

Byte	Bits	Field Name	Register Description	Type
192	7-4	-	**ACS::NPConfigLane1**	RO
	3-1	NPIDLane1	See Table 8-125	Rqd.
	0	NPInUseLane1		
193	7-4	-	**ACS::NPConfigLane2**	RO
	3-1	NPIDLane2	See Table 8-125	Rqd.
	0	NPInUseLane2		
194	7-4	-	**ACS::NPConfigLane3**	RO
	3-1	NPIDLane3	See Table 8-125	Rqd.
	0	NPInUseLane3		
195	7-4	-	**ACS::NPConfigLane4**	RO
	3-1	NPIDLane4	See Table 8-125	Rqd.
	0	NPInUseLane4		
196	7-4	-	**ACS::NPConfigLane5**	RW
	3-1	NPIDLane5	See Table 8-125	Rqd.
	0	NPInUseLane5		
197	7-4	-	**ACS::NPConfigLane6**	RW
	3-1	NPIDLane6	See Table 8-125	Rqd.
	0	NPInUseLane6		
198	7-4	-	**ACS::NPConfigLane7**	RW
	3-1	NPIDLane7	See Table 8-125	Rqd.
	0	NPInUseLane7		
199	7-4	-	**ACS::NPConfigLane8**	RW
	3-1	NPIDLane8	See Table 8-125	Rqd.
	0	NPInUseLane8		

NP States

Table 8-135 Lane-associated Network Path States (Page 16h)

Byte	Bit	Field Name	Register Description (NPStateHostLane<i>)	Type
200	7-4	NPStateHostLane2	Network Path State of host lane 2 (see Table 8-136)	RO Rqd.
	3-0	NPStateHostLane1	Network Path State of host lane 1 (see Table 8-136)	
201	7-4	NPStateHostLane4	Network Path State of host lane 4 (see Table 8-136)	RO Rqd.
	3-0	NPStateHostLane3	Network Path State of host lane 3 (see Table 8-136)	
202	7-4	NPStateHostLane6	Network Path State of host lane 6 (see Table 8-136)	RO Rqd.
	3-0	NPStateHostLane5	Network Path State of host lane 5 (see Table 8-136)	
203	7-4	NPStateHostLane8	Network Path State of host lane 8 (see Table 8-136)	RO Rqd.
	3-0	NPStateHostLane7	Network Path State of host lane 7 (see Table 8-136)	

Table 8-136 Network Path State Encoding

Encoding	State
0h	Reserved
1h	NPDeactivated (or unused lane)
2h	NPInit
3h	NPDeinit
4h	NPActivated
5h	NPTxTurnOn
6h	NPTxTurnOff
7h	NPInitialized
8h-Fh	Reserved

Table 8-137 Network Path Conditions (Page 16h)

Byte	Bits	Field Name	Register Description	Type
204	7	NPInitPendingLane8	**NPInitPendingLane<i>** 0b: NPInit not pending 1b: Commissioning the NP Active Control Set during NPInit has not yet been executed after a successful ApplyNPInit, hence the NP Active Control Set content may still deviate from the actual hardware configuration. *Note: The setting SteppedConfigOnly is irrelevant for the NPSM.*	RO Rqd.
	6	NPInitPendingLane7		
	5	NPInitPendingLane6		
	4	NPInitPendingLane5		
	3	NPInitPendingLane4		
	2	NPInitPendingLane3		
	1	NPInitPendingLane2		
	0	NPInitPendingLane1		

- Following an apply, InitPending is set to indicate the NPSM has not initialized

Replacement Signals

Page 16h allows advertising and control of replacement signals

- Register 226 advertises support for Tx and Rx path replacement signals
- Registers 162 and 163 allow selection of replacement signals or traffic

Table 8-139 Miscellaneous Options (Page 16h)

Byte	Bit	Field Name	Field Description	Type
226	7-2	-	**Reserved**	RO Rqd.
	1	ReplaceHPSignalTxSupported	0b: not supported 1b: Tx replacement signals for NP inputs are supported	
	0	ReplaceHPSignalRxSupported	0b: not supported 1b: Rx replacement signals for NP outputs are supported	

Table 8-129 Network and Host Path Signal Source Selection (Page 16h)

Byte	Bits	Field Name	Register Description	Type
162	7	HPSourceRx8	**HPSourceRx<i>** Controls the signal source feeding host lane <i> in Rx direction 0b: Signal received from Network Path 1b: Internally generated client replacement signal All host lanes belonging to the same HP must have the same HPSourceRx<i> value set Advertisement: 16h:226.0	RW Cond.
	6	HPSourceRx7		
	5	HPSourceRx6		
	4	HPSourceRx5		
	3	HPSourceRx4		
	2	HPSourceRx3		
	1	HPSourceRx2		
	0	HPSourceRx1		
163	7	NPSourceTx8	**NPSourceTx<i>** Controls the signal source feeding the Network Path input related to the HP containing host lane <i> 0b: Signal received from Host Path 1b: Internally generated client replacement signal All host lanes belonging to the same HP must have the same NPSourceTx<i> value set Advertisement: 16h:226.1	RW Cond.
	6	NPSourceTx7		
	5	NPSourceTx6		
	4	NPSourceTx5		
	3	NPSourceTx4		
	2	NPSourceTx3		
	1	NPSourceTx2		
	0	NPSourceTx1		

Multiplexing

- The NPSM can support multiplexing of mixed-rate host signals
  - Mixed rate multiplexing is accomplished by programming a set of self-consistent Appsels into the Active Set
  - All lanes must have the same Media Identifier

- Multiplexing has a restriction that All Host Lanes must be operating at the same rate:
  - 100GAUI-1, 200GAUI-2, 400GAUI-4 multiplexing can be supported
  - CAUI-4 and OTL4.4 multiplexing is NOT supported

  → *Although undesired this restriction is included to limit the advertising required for mixed rate multiplexing*

Multiplexing Advertisement

- Enumerated multiplexing structures are advertised
- Registers 228-231 advertise HostInterfaceID's indicating the lane data rate
- Registers 232-247 are a bit mask for supported Mux structures

Table 8-149 Multiplex Lane Grouping Advertisement

Multiplex Structure			HP DPID per Host Lane #							
ID	# of HPs	HP Widths	1	2	3	4	5	6	7	8
0	8	1, 1, 1, 1, 1, 1, 1, 1	1	2	3	4	5	6	7	8
1	1	8	1	1	1	1	1	1	1	1
2	2	4, 4	1	1	1	1	5	5	5	5
3	4	2, 2, 2, 2	1	1	3	3	5	5	7	7
4	3	4, 2, 2	1	1	1	1	5	5	7	7
5	4	4, 2, 1, 1	1	1	1	1	5	5	7	8
6	4	4, 1, 1, 2	1	1	1	1	5	6	7	7
7	5	4, 1, 1, 1, 1	1	1	1	1	5	6	7	8
8	3	2, 2, 4	1	1	3	3	5	5	5	5
9	4	2, 1, 1, 4	1	1	3	4	5	5	5	5
10	4	1, 1, 2, 4	1	2	3	3	5	5	5	5
11	4	1, 1, 1, 1, 4	1	2	3	4	5	5	5	5
12	5	2, 2, 2, 1, 1	1	1	3	3	5	5	7	8
13	5	2, 2, 1, 1, 2	1	1	3	3	5	6	7	7
14	5	2, 1, 1, 2, 2	1	1	3		5	5	7	7
15	5	1, 1, 2, 2, 2	1	2	3	3	5	5	7	7
16	6	2, 2, 1, 1, 1, 1	1	1	3	3	5	6	7	8
17	6	2, 1, 1, 2, 1, 1	1	1	3	4	5	5	7	8
18	6	2, 1, 1, 1, 1, 2	1	1	3	4	5	6	7	7
19	6	1, 1, 2, 2, 1, 1	1	2	3	3	5	5	7	8
20	6	1, 1, 2, 1, 1, 2	1	2	3	3	5	6	7	7
21	6	1, 1, 1, 1, 2, 2	1	2	3	4	5	5	7	7
22	6	2, 1, 1, 1, 1, 1, 1	1	1	3	4	5	6	7	8
23	7	1, 1, 2, 1, 1, 1, 1	1	2	3	3	5	6	7	8
24	7	1, 1, 1, 1, 2, 1, 1	1	2	3	4	5	5	7	8
25	7	1, 1, 1, 1, 1, 1, 2	1	2	3	4	5	6	7	7

Multiplexing Advertisement

Host Interface ID 0Dh = 100GAUI-2 : 50G lanes

Table 8-150 Multiplex Granularities Advertisement (Page 16h)

Byte	Bits	Field Name	Register Description	Type
228	7-0	MuxGranularity1	**U8 MuxGranularity<i>** 0: Not supported (end of granularity list: after a zero value, all following MuxGranularity<i> fields are zero as well) >0: HostInterfaceID indicating the lane data rate and the multiplex rate granularity: all multiplexed signals have the same lane data rate and a data rate that is a power of two multiple of the multiplex rate granularity	RO Rqd.
229	7-0	MuxGranularity2		
230	7-0	MuxGranularity3		
231	7-0	MuxGranularity4		

Table 8-151 Global Multiplex Structures Advertisement (Page 16h)

Byte	Bits	Field Name	Register Description	Type
232-235	31-0	MuxStructsSupported1	U32 **MuxStructsSupported<i>** contains a bit mask where each bit <j> set indicates support for the multiplex structure ID <j> as defined in Table 8-149	RO Rqd.
236-239	31-0	MuxStructsSupported2		
240-243	31-0	MuxStructsSupported3		
244-247	31-0	MuxStructsSupported4		

Table 8-141 Multiplex Lane Grouping Advertisement

Multiplex Structure			HP DPID per Host Lane #							
ID	# of HPs	HP Widths	1	2	3	4	5	6	7	8
0	8	1, 1, 1, 1, 1, 1, 1, 1	1	2	3	4	5	6	7	8
1	1	8	1	1	1	1	1	1	1	1
2	2	4, 4	1	1	1	1	5	5	5	5
3	4	2, 2, 2, 2	1	1	3	3	5	5	7	7
4	3	4, 2, 2	1	1	1	1	5	5	7	7
5	4	4, 2, 1, 1	1	1	1	1	5	5	7	8
6	4	4, 1, 1, 2	1	1	1	1	5	6	7	7
7	5	4, 1, 1, 1, 1	1	1	1	1	5	6	7	8
8	3	2, 2, 4	1	1	3	3	5	5	5	5
9	4	2, 1, 1, 4	1	1	3	4	5	5	5	5
10	4	1, 1, 2, 4	1	2	3	3	5	5	5	5
11	4	1, 1, 1, 1,	1	2	3	4	5	5	5	5
12	5	2, 2, 2, 1, 1	1	1	3	3	5	5	7	8
13	5	2, 2, 1, 1, 2	1	1	3	3	5	6	7	7
14	5	2, 1, 1, 2, 2	1	1	3		5	5	7	7
15	5	1, 1, 2, 2, 2	1	2	3	3	5	5	7	7
16	6	2, 2, 1, 1, 1, 1	1	1	3	3	5	6	7	8
17	6	2, 1, 1, 2, 1, 1	1	1	3	4	5	5	7	8
18	6	2, 1, 1, 1, 1, 2	1	1	3	4	5	6	7	7
19	6	1, 1, 2, 2, 1, 1	1	2	3	3	5	5	7	8
20	6	1, 1, 2, 1, 1, 2	1	2	3	3	5	6	7	7
21	6	1, 1, 1, 1, 2, 2	1	2	3	4	5	5	7	7
22	6	2, 1, 1, 1, 1, 1, 1	1	1	3	4	5	6	7	8
23	7	1, 1, 2, 1, 1, 1, 1	1	2	3	3	5	6	7	8
24	7	1, 1, 1, 1, 2, 1, 1	1	2	3	4	5	5	7	8
25	7	1, 1, 1, 1, 1, 1, 2	1	2	3	4	5	6	7	7

100GAUI-2, 200GAUI-4, 400GAUI-8 → Same Mux Granularity

Multiplexing Advertisement

4Ch: 100GAUI-1, 200GAUI-2, 400GAUI-4, 800GAUI-8

→ Same Mux Granularity : 100G Lanes

Table 8-150 Multiplex Granularities Advertisement (Page 16h)

Byte	Bits	Field Name	Register Description	Type
228	7-0	MuxGranularity1	**U8 MuxGranularity<i>** 0: Not supported (end of granularity list: after a zero value, all following MuxGranularity<i> fields are zero as well) >0: HostInterfaceID indicating the lane data rate and the multiplex rate granularity: all multiplexed signals have the same lane data rate and a data rate that is a power of two multiple of the multiplex rate granularity	RO Rqd.
229	7-0	MuxGranularity2		
230	7-0	MuxGranularity3		
231	7-0	MuxGranularity4		

Table 8-151 Global Multiplex Structures Advertisement (Page 16h)

Byte	Bits	Field Name	Register Description	Type
232-235	31-0	MuxStructsSupported1	U32 **MuxStructsSupported<i>** contains a bit mask where each bit <j> set indicates support for the multiplex structure ID <j> as defined in Table 8-149	RO Rqd.
236-239	31-0	MuxStructsSupported2		
240-243	31-0	MuxStructsSupported3		
244-247	31-0	MuxStructsSupported4		

Table 8-141 Multiplex Lane Grouping Advertisement

Multiplex Structure			HP DPID per Host Lane #							
ID	# of HPs	HP Widths	1	2	3	4	5	6	7	8
0	8	1, 1, 1, 1, 1, 1, 1, 1	1	2	3	4	5	6	7	8
1	1	8	1	1	1	1	1	1	1	1
2	2	4, 4	1	1	1	1	5	5	5	5
3	4	2, 2, 2, 2	1	1	3	3	5	5	7	7
4	3	4, 2, 2	1	1	1	1	5	5	7	7
5	4	4, 2, 1, 1	1	1	1	1	5	5	7	8
6	4	4, 1, 1, 2	1	1	1	1	5	6	7	7
7	5	4, 1, 1, 1, 1	1	1	1	1	5	6	7	8
8	3	2, 2, 4	1	1	3	3	5	5	5	5
9	4	2, 1, 1, 4	1	1	3	4	5	5	5	5
10	4	1, 1, 2, 4	1	2	3	3	5	5	5	5
11	4	1, 1, 1, 1,	1	2	3	4	5	5	5	5
12	5	2, 2, 2, 1, 1	1	1	3	3	5	5	7	8
13	5	2, 2, 1, 1, 2	1	1	3	3	5	6	7	7
14	5	2, 1, 1, 2, 2	1	1	3		5	5	7	7
15	5	1, 1, 2, 2, 2	1	2	3	3	5	5	7	7
16	6	2, 2, 1, 1, 1, 1	1	1	3	3	5	6	7	8
17	6	2, 1, 1, 2, 1, 1	1	1	3	4	5	5	7	8
18	6	2, 1, 1, 1, 1, 2	1	1	3	4	5	6	7	7
19	6	1, 1, 2, 2, 1, 1	1	2	3	3	5	5	7	8
20	6	1, 1, 2, 1, 1, 2	1	2	3	3	5	6	7	7
21	6	1, 1, 1, 1, 2, 2	1	2	3	4	5	5	7	7
22	6	2, 1, 1, 1, 1, 1, 1	1	1	3	4	5	6	7	8
23	7	1, 1, 2, 1, 1, 1, 1	1	2	3	3	5	6	7	8
24	7	1, 1, 1, 1, 2, 1, 1	1	2	3	4	5	5	7	8
25	7	1, 1, 1, 1, 1, 1, 2	1	2	3	4	5	6	7	7

Reconfiguration

- Coherent modules typically support reconfiguration of host traffic when the Media side is unchanged
  - For example, running an 800ZR media side application allows reconfiguration from 8x800GAUI-8 to 2x400GAUI-4 or (400GAUI-4)+(2x200GAUI-2) without line reconfiguration
- A Network Path requires configuration with (partial) Application Descriptors all containing the same Media ID
  - The NP Active Set is defined by the # of lanes + Host Starting Lane Number
  - This is identical for any mix of traffic with the same lane count & starting Lane
- Reconfiguration in which the Media ID Changes requires De-initializing the Network Path
  - For example, changing from 800ZR to FlexO-8-DO must first de-Initialize the NP

Threshold updates

- Prior to CMIS-5.3, no discussion of updating thresholds based on the Application was provided
  - Different Applications may however have different optical power, BER, eSNR, etc. Levels at which warnings and alarms should be raised

- In CMIS-5.3, Page 02h and VDM thresholds *may* be updated after an application change
  - For the DPSM, thresholds are updated after the DP reaches DPInitialized
  - For the NPSM
    - Media thresholds are updated when the NPSM reaches NPInitialized
    - Host thresholds are updated after the Host Side reaches DPInitialized

- Host SW should read the thresholds following (re)initialization

New topics in CMIS 5.3

- Extended Appsels (NADs) have been defined in CMIS 5.3 to expand the addressable applications – see page 1Ch.

- Host Lane Switching – Allows redirection of Host Lanes to additional data paths

Topics for evolution / further work (my opinions)

- NPSM Modules in legacy hosts
  - Legacy hosts attempting to control a module's NP application through the DPSM registers would fail
    - Defaults would currently turn on the media side of a module as soon as the Module reaches Module_Ready
- More detailed Multiplexing Applications need definition
  - CMIS restrictions on lane mixing are prohibitively and un-necessarily strict
  - Improved muxing likely advertised through a CDB command

Summary

- The NPSM was added into CMIS 5.1 to allow the media side and host side to be decoupled

- The NPSM Opens up more complex module behavior:
  - Multiplexing
  - Replacement signals
  - In service traffic changes

- CMIS is continuing to evolve to allow standardized behavior

Thank You!